

UTC DISSEMINATION TO THE REAL-TIME USER

Judah Levine
Time and Frequency Division
National Institute of Standards and Technology
Boulder, Colorado 80303

Abstract

This paper concludes the tutorial session on the delivery of UTC to the real-time user. The first paper of the session describes the technical problems and other issues that impede the ability of users to achieve synchronization and syntonization. The next four papers describe the activities of several national and international organizations that provide critical time and frequency services. These papers address many of the questions, explain the roles of the organizations and describe some of the improvements that the future may bring.

This paper summarizes the session. It synthesizes the individual presentations to provide an overall assessment of how well the real-time user can perform synchronization and syntonization using the existing infrastructure. Finally, it highlights areas where improvements can be made by providing additional information or services.

THE DEFINITION OF UTC

The current definition of Coordinated Universal Time (UTC) dates from 1972^[1]. The duration of a UTC second is defined in terms of the frequency of a hyperfine transition in the ground-state of cesium. This standard frequency is realized in a number of different laboratories using ensembles of commercial cesium clocks and a few primary frequency standards. The data from all of these devices are transmitted periodically to the Bureau International des Poids et Mesures (BIPM) in Sevres, France, where they are combined in a statistical procedure to produce International Atomic Time (TAI)^[2]. The time of this scale is adjusted as needed ("Coordinated") by adding or dropping integer seconds so as to keep it within ± 0.9 s of UT1, a time-scale based on the observations of the transit times of stars and corrected for the predicted seasonal variations in these observations. When the leap seconds are included into TAI, the result is called UTC. The difference between TAI and UTC is therefore an exact integer number of seconds. This difference is currently 29 s and will become 30 s at 0 UTC on 1 January 1996.

THE COMPUTATION OF TAI AND UTC

The algorithm used by the BIPM to compute TAI is called ALGOS, an algorithm designed to optimize the long-term stability of the average frequency of the ensemble. It is computed retrospectively using data from about 250 clocks located in many different laboratories. The computation is performed at the end of each month and the results are usually available in about 3 weeks: the computation for October, 1995, for example, was published on 16 November.

The BIPM is planning to reduce this delay, but it can never be zero because of the retrospective nature of ALGOS and because of the time needed to collect the data from the contributing laboratories. In addition, the BIPM computations use relatively infrequent measurements: timing laboratories currently report clock data to the BIPM on a 10-day mesh (at 0 UTC on every MJD ending in 9); the intervals between data points will be reduced to 5 days in the near future. The monthly computation cycle will not be changed.

Estimating UTC in real time therefore requires both extrapolation from the most recent computation of ALGOS and interpolation to a time grid more suited to real-time applications. These computations are simple in principle, but they can have significant uncertainties because the underlying noise processes in the data are not well known. Even when the noise process is known, extrapolation is an uncertain business, especially for processes characterized by flicker or random-walk spectral distributions. In addition to the uncertainties in estimating UTC in real time, the BIPM does not transmit a physical realization of UTC, so that users with real-time requirements must obtain UTC from a timing laboratory.

THE REALIZATION OF UTC BY A TIMING LABORATORY

In addition to operating ensembles of clocks, timing laboratories generally average the readings of these devices in some way to compute a local realization of UTC called UTC(lab) to distinguish it from the scale computed by the BIPM, which is written as simply UTC. As part of its monthly analysis, the BIPM computes $UTC - UTC(\text{lab})$ for each laboratory that contributes clock data and publishes these differences in a monthly bulletin called Circular T. This publication gives the value of $UTC - UTC(\text{lab})$ on the same 10-day mesh that is used to submit the clock measurements: currently at 0 UTC on every MJD that ends in 9.

The method used to realize UTC(lab) varies from laboratory to laboratory. In some cases UTC(lab) is the output of a "principal" clock whose output may be either free-running or steered towards UTC using data from Circular T. Other laboratories define a UTC scale based on a weighted average of the times of their local clocks. This is the procedure used at NIST.

The NIST clock ensemble consists of a number of commercial cesium standards and hydrogen masers. The data from these devices are combined in a time scale called AT1, which is computed automatically every 2 hours. The weight of each clock in AT1 is based on its previous performance except that no clock can have a weight greater than 30%. The algorithm is designed to average the white frequency noise that usually characterizes cesium clocks at relatively short periods of a few days or less. The scale is normally free-running – its time and frequency are not adjusted administratively. Clocks are added and dropped from the scale

as needed: a clock is added only after its performance has been evaluated for some time so as to minimize the perturbation to the ensemble average and a clock that appears to be nearing the end of its life is dropped before it actually fails.

UTC(NIST), in turn, is computed from AT1 using an equation which is designed to steer UTC(NIST) towards UTC with a time constant of several months. The parameters of the equation are estimated using the most recent 36 values (an interval of 360 days) of UTC - UTC(NIST) from Circular T. The parameters of the equation are published in the monthly NIST Time and Frequency Bulletin; each issue gives older values, the official parameters for the current month and the provisional values for one month in the future. The equation is only changed at 0 UTC on the first day of every month and includes data from the most recently-received Circular T; the equation for December, for example, is based on the BIPM computations of UTC - UTC(NIST) through the end of October. Only the rate offset is changed from month to month (time steps are never used), and the change in rate is limited administratively to be not more than ± 2 ns/day (a fractional frequency change of not more than $\pm 2.3 \times 10^{-14}$).

The time of UTC(NIST) is realized physically using a computer-controlled phase-stepper, which operates at 5 MHz. The input is from one of the clocks in the ensemble and the output is monitored every 12 minutes; these measurements are used to control the phase-stepper so as to lock the physical signal to the predicted value of UTC(NIST) - AT1. The difference between the definition of UTC(NIST) and its physical realization is about 0.2 ns RMS. This difference arises from two effects: (1) the time dispersion during the 12 minutes between measurements due to the frequency-noise of the clock driving the phase-stepper, and (2) by the white noise in the measurement process itself. The adjustments applied by the phase-stepper are typically on the order of ps, so that the frequency stability of UTC(NIST) is essentially the same as that of its parent scale AT1 for averaging times longer than the 12 minutes cycle time of the control loop.

We are experimenting with other ways of realizing UTC(NIST). In one experiment, we have added a clean-up oscillator after the phase-stepper to improve the spectral purity of the steered 5 MHz signal. This is advantageous for satellite time-transfer equipment and other systems that use the 5 MHz as a reference frequency in addition to the 1 pps output that is used by the GPS receivers. In a second experiment, we have changed the steering algorithm to emphasize frequency smoothness at the expense of time accuracy. The phase-stepper is driven from a hydrogen maser in this case, but both the amplitude and the frequency of the phase adjustments are controlled to provide maximum frequency smoothness at intermediate periods of a few days or less. The resulting output has the frequency of UTC(NIST) on the average, with almost the stability of the hydrogen maser reference. (The price of frequency smoothness is time-dispersion, and this implementation is therefore designed for users whose primary need is for frequency stability rather than time accuracy.)

In all realizations, the reference plane for the time is at a specified input to a counter located in the clock room; delays in the distribution system after that point must be measured and are included as offsets in subsequent analyses.

THE ACCURACY AND STABILITY OF UTC(lab)

Both NIST and USNO steer their respective UTC(lab) scales so that the difference UTC - UTC(lab) is small – on the order of 10 ns or less. This is not universally true, however. Many laboratories have significant offsets (both in time and in frequency) between UTC and UTC(lab). In general, therefore, the scale UTC(lab) can not be used as a replacement for UTC without estimating the difference between the times and rates of the two scales. The accuracy with which this can be done is determined by the stability of the frequency difference UTC - UTC(lab). For labs such as USNO, NIST or PTB, this value can be on the order of $(5 \pm 3) \times 10^{-15}$ for an averaging time of about 30 days, so that data from a previous Circular T can be used to predict the current value of UTC - UTC(lab) with an uncertainty of about 10 – 20 ns RMS.

SYSTEMATIC OFFSETS IN GPS DATA

Many users with demanding time or frequency requirements use data from GPS satellites to estimate UTC. These data may be used directly to estimate UTC(USNO) from the received values of GPS time and its offset to UTC(USNO). Alternatively, some of the noise in the observations may be at least partially canceled using the common-view method in which two laboratories observe a given satellite simultaneously. Common-mode errors (such as the satellite clock and a portion of the unmodeled atmospheric delay) cancel in the differences of the two measurements. (The BIPM issues tracking schedules which facilitate the simultaneous observations that are needed for common-view observations. All timing laboratories adhere to these schedules; the data from these observations form the basis for international time coordination.)

Direct and common-view observations are affected in first order by fluctuations in the transmission delay through the antenna and the receiver, by delays introduced by multi-path effects, and by fluctuations in the transit time of the reference pulse from the reference plane of the laboratory to the receiver hardware. These effects can be on the order of 50 ns or more; they are also likely to be temperature-dependent and therefore hard to measure accurately.

In addition, some laboratories use a reference for their GPS receiver that is intentionally offset from the corresponding UTC(lab) for some administrative reason. Determining these offsets accurately may be difficult because of the variations in the input impedance of the receivers which change the effective arrival time of the 1 pps pulse. Variations of this kind also change the voltage-standing-wave ratio of the cable that delivers the 1 pps; these fluctuations may be important for long cables but are quite difficult to characterize.

The BIPM has conducted several differential calibrations using a single portable receiver that is operated at each laboratory in parallel with the permanent equipment there, but the results of these comparisons are somewhat ambiguous. At some of the laboratories, the calibration constants have changed by as much as 10 ns over periods of months, while other laboratories show essentially no change over many years. The BIPM attributes some of these fluctuations to the effects of changes in the local temperature on the long cable between the antenna and the receiver, but this cause is probably not the whole story, and these effects are not completely

understood at this time^[3].

MEASUREMENT NOISE

The largest source of noise in GPS measurements is likely to be the fluctuations in time delay introduced by Selective Availability. The effect of SA can be seen in Figure 1, which shows the time difference between a local cesium clock and GPS time as measured using two satellites: SV 12 which does not have SA and SV 14 which does. Each point on the figure is a 5-minute average of the time-difference; the lines connecting the points are to make it easier to identify them and are not otherwise significant. The RMS scatter of the data from SV 14 is about 60 ns; frequency estimates computed by dividing the first difference of these data by 300 s will therefore have a scatter of about 4.2×10^{-10} .

If the frequency of the local clock is sufficiently stable, both the time and the frequency estimates can be improved by averaging. The spectrum of the GPS time-difference measurements is usually characterized by white phase noise out to averaging times on the order of 10^4 s or so; if fluctuations in the frequency of the local clock are not important on this time scale, the uncertainty in the time estimate can be reduced by about a factor of 6 by averaging for about 10^4 s. (Longer averaging times can be used as long as the residuals are still characterized by white phase fluctuations. These longer averaging times generally require additional post-processing of the data, which is not useful in a real-time context.) It is often possible to improve the estimate of the frequency offset of the local clock by additional averaging without the need for extensive post-processing; the maximum averaging time for a frequency estimator will usually be set by the time at which the fluctuations in the frequency of the local clock are no longer white. This may be several days for a good cesium standard. A number of effects may produce a variation in the time difference data with a period of 1 day, and any averaging algorithm must be designed with this in mind.

It is important to remember that averaging is only useful until the noise floor set by the characteristics of the local clock (or clock ensemble) is reached. Furthermore, while increased averaging may improve the noise performance, it degrades the transient response of the system. This means that it will take longer for the estimator to reach equilibrium initially and to recover from a glitch in steady-state operation.

CONCLUSIONS

Most real-time users who need UTC time or frequency information use the signals from the GPS satellites to get it. They can either recover UTC(USNO) using the broadcast signal directly or they can use the common-view method in which they receive signals from satellites that are being observed at the same time by a timing laboratory which realizes a local estimate of UTC.

Both of these methods of receiving UTC are limited by a number of problems. For non-authorized users, the largest problem is usually the intentional degradation imposed on the GPS signals by selective availability (SA), but systematic errors in the receiver hardware and

uncertainties in the relationship between UTC(lab) and UTC are also significant. Authorized users are not affected by SA in principle, so that the remaining effects become their dominant sources of noise.

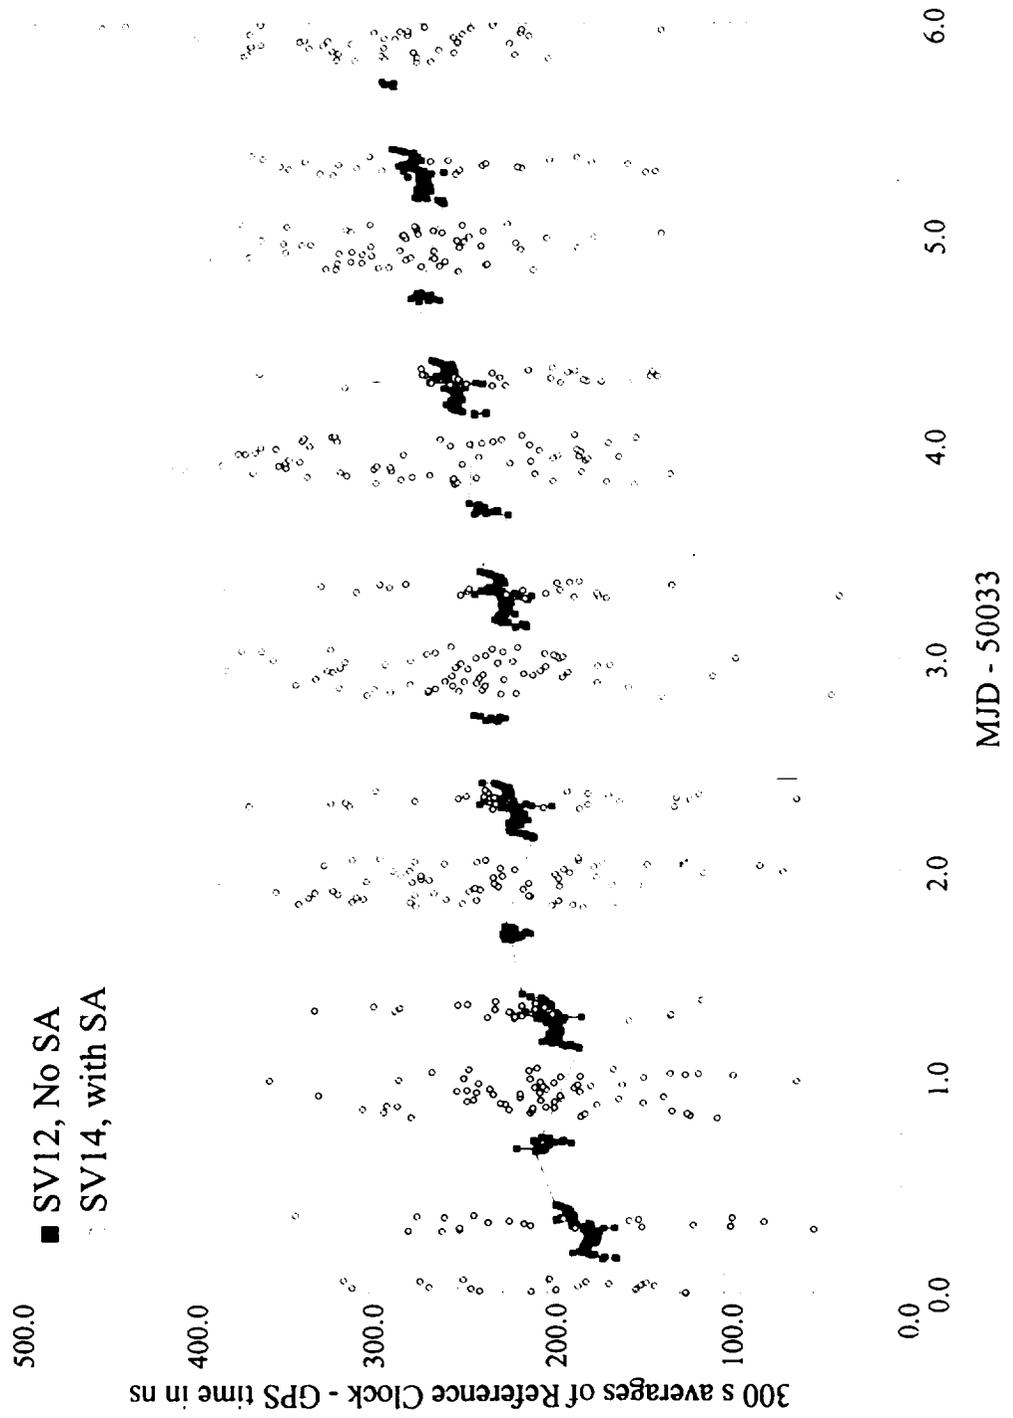
It would be nice if SA were switched off. Lacking that, it is important to design data acquisition algorithms that minimize its impact. The acquisition algorithms for GPS data that are used by all timing laboratories were designed before SA was implemented, and it is not clear that they are still optimum in this new environment. Furthermore, the increasing availability of multi-channel receivers suggests that the traditional track schedules should be augmented or modified to make better use of this hardware.

Finally, it is important for users to understand how each timing laboratory realizes its UTC(lab). In particular, it is important for users to know when time or frequency steps are applied and how large they are expected to be. These effects are probably smaller than the degradations due to SA in most cases, but this is not true for authorized users now and may not be true in the future if SA is ever turned off or if alternate distribution methods, such as two-way satellite time transfer or distribution via fibers, become widely used.

REFERENCES

- [1] T.J. Quinn 1991, "*The BIPM and the Accurate Measurement of Time*", Proceedings of the IEEE, **79**, 894-905.
- [2] B. Guinot, and C. Thomas 1989, "*The Establishment of International Atomic Time*", Annual Report of the Time Section, BIPM, Sèvres, France.
- [3] W. Lewandowski, and C. Thomas, BIPM, private communication.

Fig. 1. The effect of Selective Availability on GPS data.



Questions and Answers

CAPT. STEVEN HUTSELL (USAF): I appreciate your comments on this. You mentioned that you believe that there might be a systematic in GPS time transfer on the order of 10 nanoseconds?

JUDAH LEVINE (NIST): Yes.

CAPT. STEVEN HUTSELL (USAF): One thing that I presented in my slides earlier, we run a daily feedback loop, and theoretically – and I'd like your opinion – if there is a systematic error in time transfer by getting the daily feedback from the Naval Observatory, the Naval Observatory would be telling us how far off we are and what corrections to make. Your comment on that.

JUDAH LEVINE (NIST): Well, there are two issues here. The first one is that that loop is an authorized loop and I am a non-authorized user. Not only am I a non-authorized user, but all the data that I have access to is non-authorized data. The result of that is that it depends very critically on how you average the SA. Okay?

The second issue is that some of the data that the BIPM uses is post-processed. Therefore, there's a difference between the post-processed ephemerides and the broadcast ephemerides. Just remember that a real time user doesn't have the luxury of post-processing and has to deal with real time. And the result of that is there may be an offset between what's calculated using a post-processed ephemeris and what's calculated in real time using a broadcast ephemeris.

So it's not obvious that this problem is your problem. It may very well be my problem. It may not be in the loop between the Naval Observatory and Falcon because that's an authorized loop.

DR. GERNOT WINKLER (USNO, RETIRED): That's a very interesting discussion. I have to agree that there are systematic problems in the absolute delay in GPS receivers; and that there are some things which we do not quite understand; and I would agree that probably the level of uncertainty is on the order of 10 nanoseconds.

But the way in which you arrived to that conclusion is completely wrong. Because, you don't see UTC(USNO). You don't get it. What you get is mean USNO, and that's not being steered. And Mr. Miranian explained this morning, we have the difference between the unsteered time scale and the steered time scale. The steered time scale, you don't see. The only instrument which sees that is the servo-loop that controls the Master Clock. So you cannot base your conclusions on that.

What you could base it on, however, is something which is very inconvenient. This is a difference if you compare in the Circular T the values given by the Paris Observatory and the Naval Observatory, they would find a discrepancy. That discrepancy is entirely viewed to the difficulty of guaranteeing, in an absolute sense, the calibration of the GPS delays.

That is a problem, and I think we have to address that sooner or later. At the moment, however, nobody is worried about it because everything is relative and you calibrate these receivers side-by-side. But, the absolute value is uncertain, and we have never actually been able to have all the receivers in one place. That would solve the problem.

JUDAH LEVINE (NIST): To answer your question, the data – when I said UTC(USNO), that came from Circular T; that's a Circular T number. I cannot see inside your servo-loops, I just looked at the Circular T.

DR. GERNOT WINKLER (USNO, RETIRED): But, it is not UTC(USNO).

JUDAH LEVINE (NIST): It says “UTC, USNO” on the piece of paper.

DR. GERNOT WINKLER (USNO, RETIRED) It also says A.1. A.1 is the unsteered time scale.

CLAUDINE THOMAS (BIPM): On Circular T, there is a note about that. What we call “UTC(USNO)” is, in fact, the USNO Master Clock. I think it's on the equation, which was wrong. What you wanted to say is that there is a difference between taking the raw data from two labs, computing the common-view observation, and the results which are in Circular T. That's what you - - -

JUDAH LEVINE (NIST): No. When I get a file from your computer, or when I get a file from the USNO computer, it says on the top of the heading “Master Clock minus GPS.” When I take that number, I take that number as Master Clock minus GPS. And when I said “USNO Master Clock,” that's where the number came from. It came from that column.

DR. GERNOT WINKLER (USNO, RETIRED): The servo-control of the GPS timing is derived from a different receiver, as compared to that which are publically available. They have to be. We have to keep that separate. It's a completely separate receiver, and there is also a question about how constant these are. There isn't a moment that this is going on between a GPS receiver, which is on loan from the BIPM, and the 502 reference receiver that we have at the Observatory, and there is no large variation.

Mr. Miranian, you have conducted that test between the two receivers. What was the result of that?

MIHRAN MIRANIAN (USNO): We've had the two receivers running side-by-side since July. And we see no more than a variation of a couple of nanoseconds. I think I showed you this last night, Judah.

And we've had temperature changes of up to 40 degrees. So this is one of the things we were trying to check on.

JUDAH LEVINE (NIST): Let me make clear what I did, okay? When I say “USNO Master Clock,” that comes from the data which you put on your computer, which the Naval Observatory puts on its computer. It says on a top of a column “Master Clock minus GPS.” Then I took that data and went around to the various laboratories which published the same data and took the differences. And then I compared that number against the values that are published by Dr. Thomas in Circular T. That's where that difference comes from.

CLAUDINE THOMAS (BIPM): Yes, because the time comparisons, which are computed at the BIPM, the computation is chosen so that to get the most precise result, but also the most accurate result. So we have already shown that to improve accuracy, we need to use, for instance, measurements, post-processed, precise ephemerides and so on. And we also add some

differential delays between some receivers, which have been obtained through experimental calibration trips. And that's why you get those differences, for sure.

WLODZIMIERZ LEWANDOWSKI (BIPM): I'm working with Mihran Miranian on this comparison between USNO and the Paris Observatory. Just a comment about the differences that we have found and what is going on.

It happens that in June '94 we compared USNO receivers to the ones we have at the Paris Observatory. We found a 13 nanosecond difference, which is not the same as the correction we just added in Circular T. This is an experiment. In Circular T, I believe there is like 20 nanoseconds now. We found 13 nanoseconds and a very clear indication that USNO receivers have a problem with temperature. We even made the determination of the coefficient which was a half of nanosecond per degree Celsius.

Then we repeated the comparison in March of 1995. And the comparison was then 20 nanoseconds between the Paris Observatory and USNO, seven nanoseconds off the previous review, and this matched exactly with the coefficient of temperature.

Then we shifted the receiver at the beginning of August, and theoretically we should receive 10 nanoseconds, according to this coefficient, and we received 10 nanoseconds. The trouble is that after several months, this is still 10 nanoseconds, and it should go down. It was going up. So there is a problem. Yesterday Dr. Gifford showed us some data which shows that the receiver jumped in behavior from one year to another, it changed behavior. So there is some work that needs to be continued concerning the calibration of receivers.